Center Overview

http://ncr.mae.ufl.edu/aacoe.php

UF UNIVERSITY of FLORIDA Duke UNIVERSITY TEXAS The University of Texas at Austin UC SANTA CRUZ

Center Overview

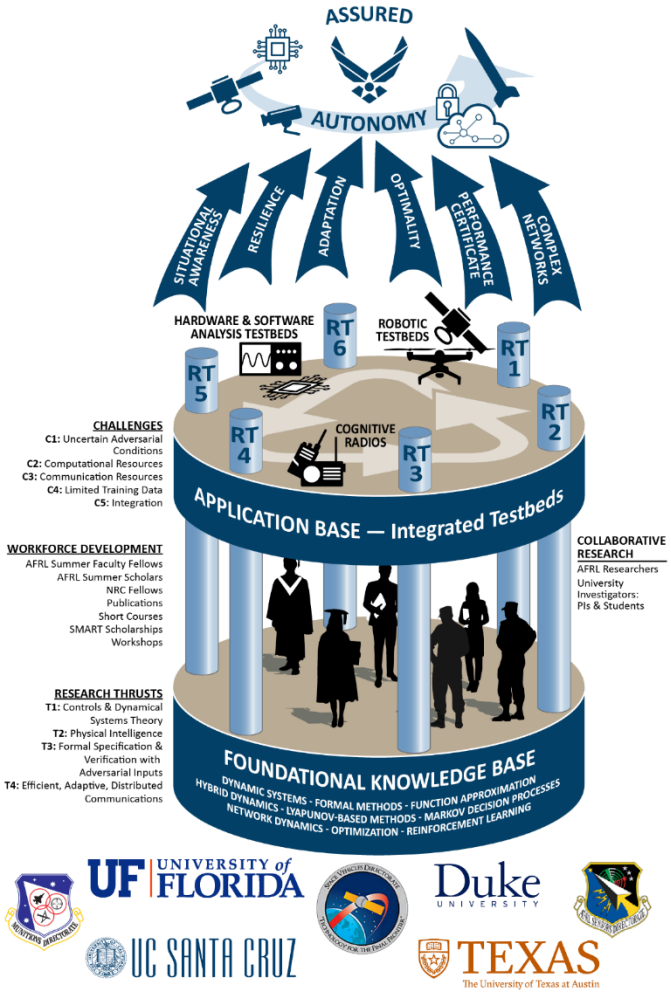

AFOSR Center of Excellence in Assured Autonomy in Contested Environments

- >$6M over 6 years (3 x 2 year increments)
- 10 PIs @ 4 Universities:
  - R. Bevilacqua (UF: optimal, switching)
  - K. Butler (UF: cyber resiliency/privacy)
  - W. Dixon (UF: ADP, networks, hybrid)
  - N. Fitz-Coy (UF: optimal, games)
  - M. Hale (UF: networks, privacy)
  - M. Pajic (Duke: cyber resiliency/privacy)
  - R. Sanfelice (UCSC: hybrid, networks)
  - J. Shea (UF: networks, privacy)
  - U. Topcu (UT: formal, hybrid, optimal)
  - M. Zavlanos (Duke: ADP, networks, formal)
- AFOSR provides 50% of funding
- AFRL (RV, RW, RY) provide 50%

UF UNIVERSITY of FLORIDA Duke UNIVERSITY TEXAS The University of Texas at Austin UC SANTA CRUZ

Center Motivation

- Innovation & technology dominance and strong economy have allowed for exquisite systems that for decades have operated in largely uncontested environments
  - Remote piloted vehicles (RPV) and monolithic satellites provide various strategic and tactical advantages
  - Intelligence, surveillance, and reconnaissance (ISR) in close proximity with RPVs or from protected space assets, while simultaneously striking from distances and with speeds beyond the capability of countermeasures
- These advantages are mitigated as the technology gap closes and as other world economies become near peers and risks to the warfighter and financial costs increase and tactical capabilities become stressed when military operations are in contested or denied environments (i.e., anti-access/area denial (A2AD) environments)
- Increased stand-off distance, persistence, and scaled projection of power have resulted in an urgency for development and fielding of human-in-the loop/semiautonomous systems

Center Motivation

- As these advantages are taken to the limit, coupled with the resultant need for rapid decision-making capabilities, emerging technology will move along a spectrum towards greater automation with less human intervention
- In contested environments, autonomous systems are even further motivated by the potential desire to complete mission execution when communication with a human operator is unavailable
- Autonomous systems must execute high level missions plans with verifiable assurances despite uncertain adversarial environments where the integrity and availability of sensor information and communications are challenged
- Key innovations include analysis, design and synthesis tools that enable autonomous mission execution despite uncertainty within complex dynamics while accounting for the integrity and privacy of information on computationally constrained resources

Center Goals & Vision

- Networks of autonomous systems will require information exchanges of many data types, including high-level mission specifications and sensor feedback for navigation and control
- The goal of assuring autonomy is complicated by the interplay between dynamics of autonomous agents and the stochastic and intermittent dynamics of network traffic
- This challenge is further amplified by delays and asynchrony in information flows
- Information perturbations can also emanate from adversarial actors in unique and complex ways, requiring security-aware design and analysis methods
- For example, we will develop techniques to protect mission-critical information and prevent information disruption/corruption
- These challenges must be addressed considering resource limitations and quantitative tradeoffs.

Research Topics

- Nonsmooth Systems
- Adaptation, Optimality, and Synthesis
- Network Systems
- Asynchronous Information
- Attack-Resilient Design
- Protecting Information

UF UNIVERSITY of FLORIDA · Duke UNIVERSITY · TEXAS The University of Texas at Austin · UC SANTA CRUZ

Workforce Dev. AFRL Collaborations Publications

UF UNIVERSITY of FLORIDA

Duke UNIVERSITY

TEXAS The University of Texas at Austin

UC SANTA CRUZ

Collaborative Interactions

- Project partially supports
  - 10 (+5) postdocs/research scientists, >40 PhD

- Now have Alumni
  - 5 postdocs – Univ. of Sherbrooke, Univ. of Arizona, Apple, Univ. Grenoble Alpes, UC Berkeley
  - 7 PhD – RW (x2), Ford, Qualcomm, Intel, Univ. of the Bio, Opener
  - 2 MS – Lockheed Martin (Orlando), Walmart Labs

- SMART Fellow for RV (Sage Edwards)

- AFRL/Space Scholar/intern
  - Summer 2021 (RY/RV (Topcu), RY/Act3 (Butler), RY (Hale), RW (Dixon)

- AFRL Summer Faculty Fellows program
  - Riccardo Bevilacqua (2019 & 2020 RW, 2021 RV)
  - Matthew Hale (2020 RW)

UF UNIVERSITY of FLORIDA Duke UNIVERSITY TEXAS The University of Texas at Austin UC SANTA CRUZ

Collaborative Interactions

- Publications
  - 196 total, ~70 this reporting cycle
  - Joint publications - 29 w/ PIs, 21 w/ AFRL

- Other
  - SIAM workshop
  - Keynote, plenary, and invited talks (DoD Basic Research Forum, Feb.)
  - 2 certificate programs proposed at UF (Astronautics, Autonomy)
  - Started bi-monthly seminar series

- Testbed Development
  - Purchase orders submitted
  - Finalizing permitting and hope to break ground in May
  - International collaborations in discussion (Brazil)
  - Certification courses and other educational outcomes

- AFRL Testbed Interactions
  - Collaborating with Chris Peterson (RV)

- Seeking Industrial Outlets for IP

Specific Research Thrusts and Advancements

UF UNIVERSITY of FLORIDA

Duke UNIVERSITY

TEXAS The University of Texas at Austin

UC SANTA CRUZ

Tightly Coupled RTs

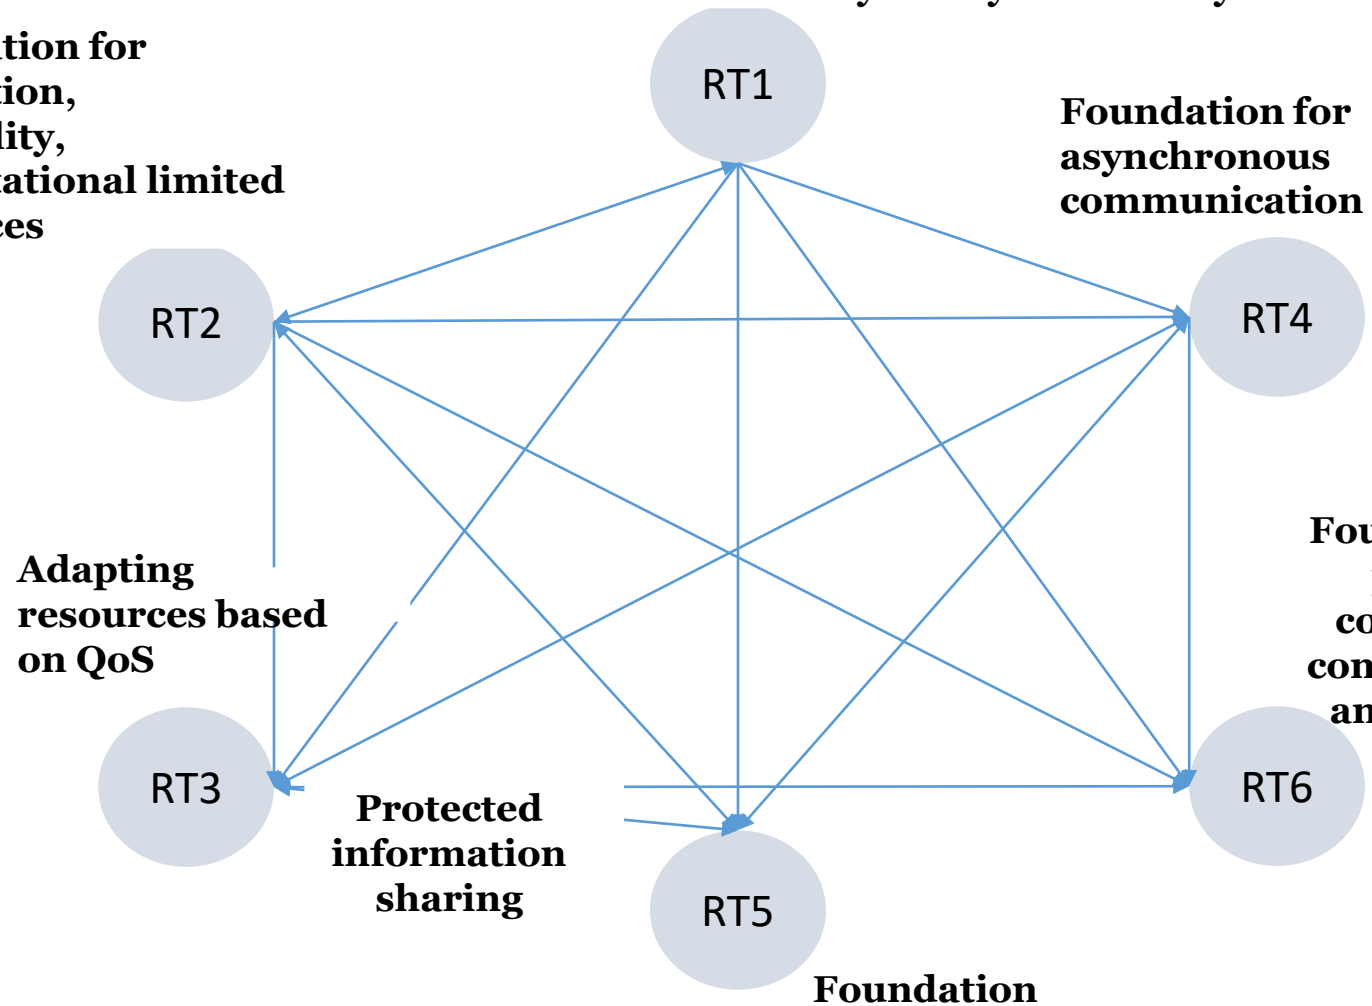

Nonsmooth Systems

- Most of the results in the center involve complex systems that have multiple modes of operation, interaction with multiple agents, sequent of events that trigger different actions, decision logic that must be reasoned about, etc. – the fusion of information and dynamics
- This RT permeates most of the other RTs, but open questions remain regarding the development of new nonsmooth analysis and design frameworks for safety, intermittency, asynchrony, uncertainty

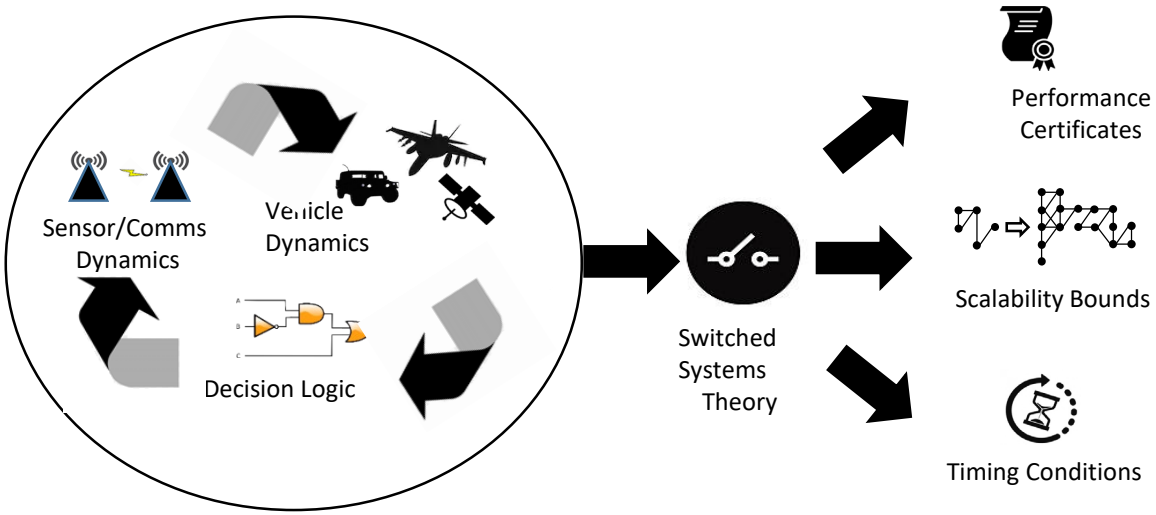

Nonsmooth Systems

- **Some Major Contributions**
  - Infinitesimal conditions for forward invariance and contractivity of sets for hybrid dynamical systems using multiple barrier functions, with applications for safety
  - Topologically motivated methodology for extending relay-explorer problems to near-arbitrary feedback region geometries, in all dimensions, with significant improvements of maximum dwell-time bounds
  - Adaptive hybrid approach to clock-synchronization over a network of agents
  - Introduced techniques (and a tool) for the probabilistic conformance of cyber-physical systems, based on a new notion of (approximate) probabilistic conformance for sets of complex specifications expressed using Signal Temporal Logic (STL)
  - LMI-based observer with intersample injection to exponentially estimate, with a given decay rate, the state of a continuous-time Lipschitz nonlinear system in the presence of sporadically available measurements.

Adaptive, Optimal, Synthesis Methods

- Complex environments often exhibit uncertainty, yet, there is a desire for autonomous systems to make the best possible (optimal) decision despite the uncertainty, and in a timely manner e.g., to meet high-level specifications
- Further complexities emerge when the system is changing among different subsystems, or the uncertainty is time-varying

- **Some Major Contributions**
  - (AFRL co-authored) Deep neural networks (DNN) for online approximation of drift dynamics, which are updated using a Lyapunov-based policy that facilitates asymptotic convergence
  - (AFRL co-authored) Used DNN system identification approach to solve an infinite horizon approximate dynamic programming problem for a system with completely unknown drift dynamics
  - Introduced a model-free learning method for design of control strategies that maximize the worst-case probability that a desired linear temporal logic (LTL) objective is satisfied in unknown non-deterministic environments

Adaptive, Optimal, Synthesis Methods

- **Some Major Contributions** (cont.)
  - (AFRL co-authored) A multi-agent reinforcement learning-based controller was developed that leverages homogeneous agent properties to facilitate better value function approximation
  - (AFRL Co-authored) Developed a data-driven state estimator for a monocular camera to estimate the Euclidean distance to features on a stationary object and to estimate the Euclidean trajectory taken by the camera while tracking the object
  - Developed synthesis algorithms for uncertain partially observable Markov decision processes that offer three-order-of-magnitude better scalability compared to the previous state of the art
  - Developed a new sampling-based LTL planning algorithm that does not require any discrete abstraction of robot mobility, is probabilistically complete and asymptotically optimal, and scales to much larger systems (number of robots, complexity of environments and tasks) compared to state-of-the-art methods

Network Systems & Asynchronous Information

- Ensuring the availability and integrity of information in complex contested environments presents a number of challenges for autonomous agents
- These challenges are exacerbated for large networks of agents operating in adversarial conditions (e.g., issues of scalability, computational resources, verifiability, collaboration strategies, and additional unique complexities)
- Since such networks are the union between the physical agent dynamics and the information dynamics, this research topic focuses on analysis, design and synthesis methods for agents *within a network and over a network*
- Moreover, sharing information across the network can lead to variable and unknown time-delays and asynchrony among agents in the network

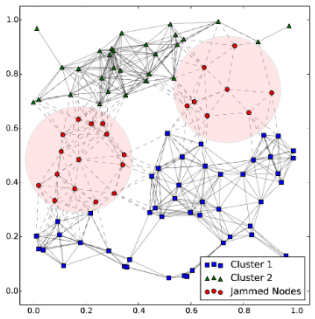

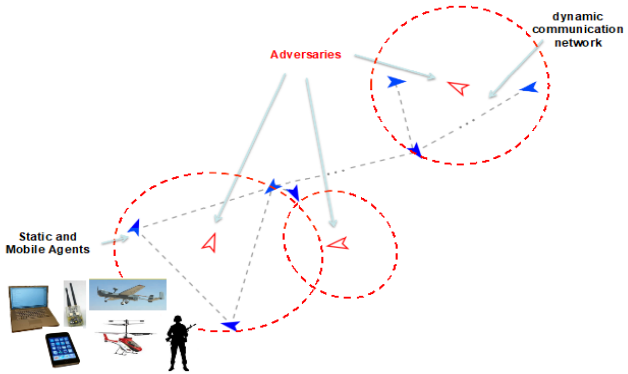

Network Systems & Asynchronous Information

- **Some Major Contributions**
  - (AFRL co-authored) Control scheme enabling clusters of agents to obtain consensus despite intermittent and asynchronous communication between clusters over a connected union graph
  - Distributed controller enabling a multi-agent system to explore an unknown, feedback-denied environment indefinitely
  - Developed bounds on the probability of connected random graphs and on the connectedness of unions of random graphs
  - Distributed multi agent optimization for networks with multiple objectives
  - Distributed online learning with constraints that vary arbitrarily over time
  - Developed distributed synthesis algorithms for with hundreds of heterogenous systems each of which with 100s of thousands of states
  - Developed new methods to transfer experience from a demonstrator agent that has access to contextual information, to a learner agent that does not have access to the environmental context, so that the learner can learn a control policy using much fewer samples

Attack Resilient Design

- The goal of this research topic is to address security challenges related to attacks against a multi-agent mission where the attacker can (1) take over a sensor and supply wrong or untimely sensor readings, (2) disrupt actuation, (3) affect communication between agents, or (4) even compromise some of the agents involved in the mission
- These attacks manifest themselves as malicious interference signals, and the defenses against them have to be introduced in the control/autonomy design
- Our key research results have focused on security-aware control design for systems with varying levels of autonomy, while guaranteeing the desired levels of system performance (i.e., Quality-of-Control) even when the system is under attack. Our goal has been to add resiliency at every level of the autonomy-stack (low-level control, switched controllers, planners)

Attack Resilient Design

- **Some Major Contributions**
  - Development of trust and reputation model using redundant state information to detect and exclude Byzantine adversaries from a network composed of cooperative agents
  - Control scheme that enables a network of agents to detect and mitigate false data injection attacks
  - Showed that a Deep Q-Network trained to perform dynamic spectrum sharing in a mobile radio scenario can improve throughput across teams of networks by careful reuse of spectrum across space
  - New optimal graph-search algorithm identifying malicious attacks and securely estimating the states even in large-scale CyberPhysical Systems
  - Introduced a fully model-free reinforcement learning method to derive security-aware control policies for unknown stochastic environments, such that even in the presence of attacks on control signals (i.e., actuators)
  - Investigated the problem of statistical model checking (SMC) for hyperproperties in security and privacy applications, focusing mainly on information security in various applications

Protecting Safety- and Mission-Critical Information

- Agents' actions can be observed by adversaries and decisions to change modes, e.g., from surveying an area to pursuing a target, can reveal modes of operation and switching strategies to adversaries
- Agents' communications can also be intercepted, including communications within and between teams of agents, and these communications can contain sensitive information, such as the agents' intents and tactics
- While adversarial environments can impose a variety of disturbances from external sources, efforts in this research topic will also explore the strategic disruption of information as a means to safeguard agents' actions and communications

Protecting Safety- and Mission-Critical Information

- **Some Major Contributions**
  - Developed a differential privacy mechanism to protect characteristics of MDPs at runtime
  - Formation control of a network of agents with differentially private exchange of information
  - Formalized deception in supervisory control and showed that while synthesis of deceptive policies is theoretically a manageable problem, synthesis of supervisory policies that protect against deception is hard
  - Adversial ML analysis of ML-based device authentication systems was carried out, demonstrating systematic weaknesses due to brittleness of extracted features in some systems, and analyzing the robustness of others
  - A framework mitigating privacy risks in augmented and virtual reality systems was developed to greatly reduce identification rates while maintaining less than 1.5 degrees of error of gaze prediction